

What Differentiates a Winning Agent: An Information Gain Based Analysis of TAC-SCM

James Andrews, Michael Benisch, Alberto Sardinha and Norman Sadeh

School of Computer Science

Carnegie Mellon University

{jandrews, mbenisch, alberto, sadeh}@cs.cmu.edu

Abstract

The Supply Chain Trading Agent Competition (TAC SCM) was designed to explore approaches to dynamic supply chain trading. During the course of each year's competition historical data is logged describing more than 800 games played by different agents from around the world. In this paper, we present analysis that is focused on determining which features of agent behavior, such as average lead time or selling price, tend to differentiate agents that win from those that don't. We begin with a visual inspection of games from one bracket of the 2006 semi-final rounds. Plots from these games allow us to isolate behavioral features which do, in fact, distinguish top performing agents in this bracket. We introduce an information gain based metric that we use to provide a more complete analysis of all the games from the 2006 quarter-final, semi-final and final rounds. The technique involves calculating the amount of information gained about an agent's performance by knowing its value for each of 20 different features. Our analysis helps identify features that differentiated winning agents. In particular we find that, in the final rounds of the 2006 competition, winning agents distinguished themselves by their procurement decisions, rather than their customer bidding decisions. We also discuss how the information gain analysis could be extended by agent developers to identify potential weaknesses in their entry.

Introduction

As the Internet helps mediate an increasing number of supply chain transactions, there is a growing interest in investigating the potential benefits of more dynamic supply chain practices (Arunachalam & Sadeh 2005; Sadeh *et al.* 1999). Since its inception, the Supply Chain Trading Agent Competition (TAC SCM) has served as a competitive test bed for this purpose (Collins *et al.* 2006). TAC SCM pits against one another trading agents developed by teams from around the world, with each agent using its own unique strategy. Agents are responsible for running the procurement, planning and bidding operations of a PC assembly company, while competing with others for both customer orders and supplies under varying market conditions.

During the course of each year's competition more than 800 games are played. The logs of these games provide ample data for evaluating the strengths and weaknesses of

more adaptive strategies than traditional supply chain management techniques (such as those described in (Chopra & Meindl 2004)). The primary measure of an agent's performance in TAC SCM is its average overall profit. Using this metric we are able to determine which agents perform best across a wide variety of conditions. However, examining only average profit does not tell us what differentiated winning agents from the others. Answering this question is of practical interest to agent designers, and may also help transfer insights from the competition to real world supply chain problems. In this paper, we investigate which features of agent behavior were most able to distinguish top performing TAC SCM agents in data from the 2006 competition's quarter-final, semi-final and final rounds.

We begin with a close look at statistical features of 6 different agents in one bracket of the 2006 semi-final rounds, such as the average quantity they requested on component orders each day. Plots from these games reveal unique patterns, or "fingerprints," which allow us to isolate behavioral features that distinguish top performing agents in this bracket.

Then, using a quantitative analysis technique, we estimate the ability of 20 different features to differentiate winners over a large collection of games. Our technique involves calculating the amount of information gained about an agent's performance by knowing its value for each feature. Our results on data from the 2006 final rounds include a ranking of features based on their information gain, providing insight into the collection of features that made winning agents unique. In particular we find that, in the final rounds of the 2006 competition, winning agents distinguished themselves by decisions related to the procurement of components, rather than those related to bidding for customers.

The remainder of this paper is organized as follows: We first provide a brief overview of the TAC SCM game. The next section describes related efforts in analyzing and presenting tools to analyze the TAC SCM games. Section 4 describes our visual inspection of feature plots from the 2006 semi-finals. In Section 5 we present our information gain-based analysis and apply it to all of the games from the 2006 quarter-finals, semi-finals and finals. The final section discusses additional uses of our technique such as how it could be extended by agent developers to identify potential weaknesses in their entry.

TAC SCM Overview

The TAC SCM game is composed of 220 game days, and on each day agents are required to make several decisions. They are responsible for sending requests to suppliers, offers to customers, and a production plan to their factory. Each request to a supplier for a specific component includes a quantity, lead time (the number of days before the order is delivered), and reserve price (the maximum the agent is willing to pay for the parts). Suppliers then respond with offers that specify actual quantities, prices, and lead times. When an agent places an order, parts are scheduled for delivery into their stock. Agents can also respond to requests made by customers for finished PCs. These requests specify a quantity, due date, reserve price, and PC type. Agents compete for each customer request by submitting offers with a specific price. The agent with the lowest price for each request is awarded an order and upon delivery the revenue for the transaction is placed in its bank account. For a more detailed description of the game, readers are directed to (Collins *et al.* 2006).

Related Work

Several researchers in the Trading Agent Competition community have presented methods for analyzing competition data to gain insights about agent performance.

In (Kiekintveld, Vorobeychik, & Wellman 2006) and (Wellman *et al.* 2006) the University of Michigan team applied game theoretic analysis to abstracted versions of the TAC games. The abstracted games were estimated empirically from the results of repeated simulations with different combinations of strategies. Their analysis revealed interesting best response and equilibrium relationships. The Michigan team also presented methods for estimating the efficiency and power of different entities in the TAC SCM market (Jordan *et al.* 2006).

In (Benisch *et al.* 2006) we analyzed data from the seeding rounds of the 2005 competition to determine that the strong performance of our agent, CMieux, was largely attributable to significantly cheaper component purchase prices than other agents.

Toolkits such as our Analysis Instrumentation Toolkit (Benisch *et al.* 2005) and the Swedish Institute for Computer Science (SICS) Game Data Toolkit¹ allow teams to analyze historical log files from a single TAC SCM game. These tools provide an in-depth view of the B2B and B2C interactions through graphical front-ends.

Several teams have also analyzed controlled experiments using different configurations of their own agent and publicly available agent binaries.

In (Borghetti *et al.* 2006) the team from the University of Minnesota presented techniques to manipulate the market environment of the simulator. By controlling various market factors, such as aggregate demand and supply, they suggest that TacTex, a top performing agent, loses its edge when market pressure is high. In (He *et al.* 2005) the Southampton team presented experiments with variants of their own agent that are more or less risk seeking in choosing selling

prices, and in (He *et al.* 2006) they provide similar analysis with respect to lead times on component orders. In (Pardoe & Stone 2006) the University of Texas team evaluated variants of their own agent against publicly available binaries of other agents. They used the results of their experiments to fine-tune various parameters in their final agent and guide future development.

The analysis methods presented in this paper differ from existing techniques in the following two ways: i.) we systematically investigate the question of which behavioral features are associated with successful performance across all agents in the 2006 final rounds and ii.) we perform all of our analysis on actual competition data, as opposed to offline controlled experiments.

Feature Plot Analysis

Historical data from the TAC SCM competition provides a large data source for studying the effectiveness of different supply chain trading techniques. However, note that by analyzing historical data we are limiting ourselves to considering only low-level actions taken by each agent, since this data does not describe underlying algorithms or techniques. In this section, we analyze plots of statistical features of these actions for six different agents from the 2006 semi-finals games containing our agent, CMieux. The data set consists of 16 games with agents placing in the following order: DeepMaize, Maxon, Botticelli, CMieux, Mertacor, and SouthamptonSCM.

Out of all the feature plots we examined, the following best illustrate how agents can be distinguished by features of their low-level behavior. Each of the plots presented shows qualitative differences between the six agents. By analyzing these plots we are able to identify unique characteristics of the agents, and gain insights into why some performed better than others.

Lead time vs. Game day

Figure 1 shows plots of the average component order lead time (Y axis) on each game day (X axis) of the different agents². These plots show that agents are easily distinguished by the extent to which they used long lead times early in the game, the length of their maximum lead time, and their most commonly used lead times.

The two best performing-agents from this round, DeepMaize and Maxon, feature substantially longer early-game lead times. They also both reduce their lead times well before Mertacor, CMieux, and Botticelli. The latter three appear to maintain long lead times until absolutely necessary. SouthamptonSCM takes a hybrid of these two approaches, reducing lead times before necessary but still much later in the game.

Maxon and Mertacor take very different approaches to the mid-game, with Mertacor almost exclusively using longer lead times, and Maxon primarily relying on short ones.

²Plots presented in this section examine behavior with respect to one specific component. Aggregating data across multiple components washed out potentially interesting details, and plots for other components were not noticeably different.

¹Available at <http://www.sics.se/tac/>.

Maxon also seems to exhibit a single mid-game 'spike' in lead times, placing orders with uncharacteristically long lead times near day 120. This is either a fixed restock point or an attempt to disrupt the procurement of other agents. Mertacor's, and, to a lesser extent, SouthamptonSCM's plots show 'bands,' which most likely correspond to specific long-term order lead times that are chosen to simplify their decision processes.

Lead time vs. Order quantity

Figure 2 shows plots of the average lead time of component orders (Y axis) against their average quantity (X axis). These plots illustrate that agents differ in the extent to which they place large orders with long lead times.

Placing component orders with long lead times and large quantities corresponds to increased risk. Thus, the extent to which an agent is willing to increase both can be seen as a reflection of its attitude towards risk. The lead time vs. order quantity plots showcase the different approaches of the agents: Maxon, Mertacor, CMieux and DeepMaize each appear reluctant to place orders with long lead times and large quantities. The trade-off is less pronounced for Botticelli and SouthamptonSCM. Maxon, Mertacor and DeepMaize each show unique 'bands,' with DeepMaize considering only a handful of fixed order quantities, Mertacor considering only fixed lead times, and Maxon fixing a combination of the two attributes.

Reserve price vs. Order price

Figure 9 in the Appendix shows a plot of each agent's average component order price (Y axis) against that agent's average offered reserve price (X axis). This plot illustrates that agents employed variations of three different strategies for choosing their reserve prices: fixed reserve prices, dynamic reserve prices, and reserve prices equal to purchase prices.

Maxon and Mertacor appear to choose from a few fixed reserve prices. SouthamptonSCM and Botticelli appear to use their reserve prices to more aggressively limit their order prices, since they are consistently close to their purchase prices. CMieux and DeepMaize have more dynamic strategies for choosing reserve prices, although a few 'bands' of fixed reserve price do appear in the DeepMaize plot.

Order quantity vs. Game day

Figure 10 in the Appendix shows a plot of each agent's average order quantity (Y axis) on each game day (X axis). Agents demonstrate unique choices for maximum order quantity, minimum order quantity, and the specific quantities they ordered repeatedly.

Mertacor, DeepMaize, and, to a lesser extent, Maxon each appear to favor orders greater than roughly 100 components at the beginning of the game. Maxon chose a maximum order quantity of about 200 units after the beginning of the game, while SouthamptonSCM and CMieux appear to consider at most about 400. Botticelli, Mertacor, and DeepMaize are all willing to go above 800 units on occasion. Bands on the graphs of DeepMaize and Southamp-

tonSCM suggest these agents were frequently choosing the same quantity on their orders.

Order price advantage vs. Lead time

Figure 11 in the Appendix shows a plot of each agent's average order lead time (Y axis) against the average order price "advantage," or the difference between their price and the best price (X axis). In these plots, agents can be distinguished by the extent to which they require better price advantages to consider long lead times.

Maxon and DeepMaize, for example, have a clear 'triangle' structure to their graphs, implying that they were only willing to accept orders with long lead times when they could get them at relatively good prices. Mertacor, SouthamptonSCM and Botticelli's plots have almost rectangular shapes, implying a more general acceptance of long lead times. CMieux appears to have a hybrid approach, with the triangle structure only being apparent for lead times above about 25 days.

Information Gain Analysis

In order to extend our analysis to a larger data set, we operationalize our notion of agent differentiation with a quantitative technique. Our technique considers the correspondence of particular features with top performance, or their *information gain*, and provides insight into the collection of features that made winning agents unique. By using a metric for comparing several different features at once, we are able to rank more than 20 different features across all 80 games from the 2006 final rounds.

Measuring information gain

In this analysis we calculate the amount of information gained about an agent's performance by knowing its value for different features. Information gain is a popular measure of association in data mining applications. The information gained about an outcome O from an attribute A is defined as the expected decrease in entropy of O conditioned on A . The following equations can be used to calculate the information gained about a discrete outcome O from a discrete attribute A , which we denote as $IG(O, A)$. We use $H(O)$ to denote the entropy of O , $H(O | A)$ to denote the entropy of O given A , and $P(a)$ to denote the probability that attribute A takes on value a in the data.

$$IG(O, A) = H(O) - H(O | A)$$

$$H(O) = - \sum_{o \in O} P(o) \log_2(P(o))$$

$$H(O | A) = \sum_{a \in A} P(a) H(O | A = a)$$

$$H(O | A = a) = - \sum_{o \in O} P(o | a) \log_2(P(o | a))$$

Intuitively, $IG(O, A)$ is how much better the value of O can be predicted by knowing the value of A . For a more detailed explanation of information gain as used in this paper see, for example, (Mitchell 1997) pp 57–60.

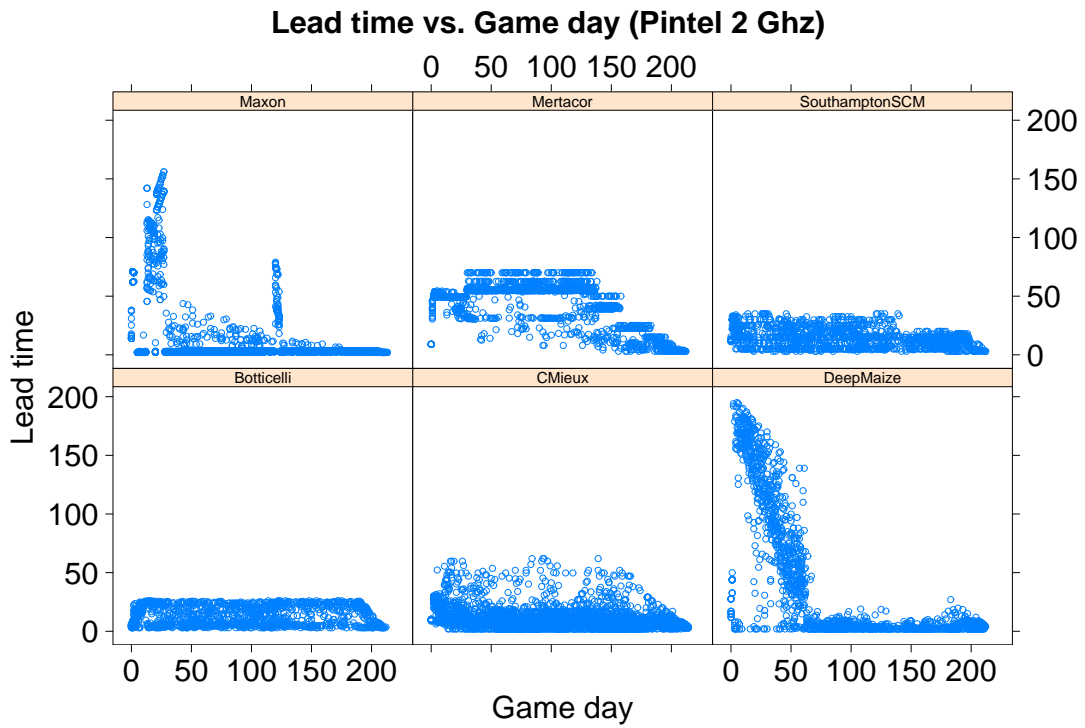


Figure 1: A plot showing the average lead time and game day of component orders placed by six different agents during the 2006 semi-finals.

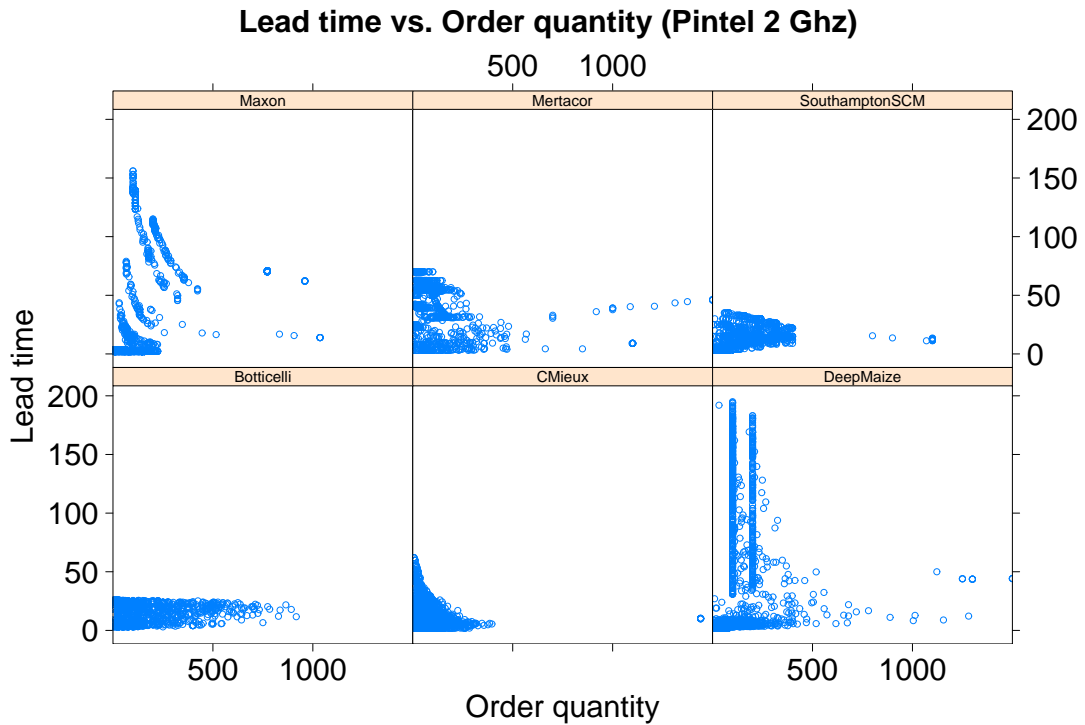


Figure 2: A plot showing the average lead time and average order quantity per day of component orders placed by six different agents during the 2006 semi-finals.

In our analysis we use the information gain metric to determine how much better we can predict an agent’s success by knowing features of its behavior. For our dataset, we construct a collection of *performance* observations, with one observation for each agent in each game. Performance observations include an outcome value, indicating whether or not the agent placed first³ and 20 different real-valued attributes of its behavior.

Before we can calculate the information gain of the attributes, we must discretize them. This is accomplished by splitting the space between the minimum and maximum values of each attribute evenly into 2^k partitions, for a positive integer k . In our results we present the information gain of all different features with k varied between 1 and 6. For a particular attribute, using larger values of k will tend to increase (and *cannot* decrease) its information gain⁴. Therefore, using values of k that are too large can lead to a kind of “over-fitting,” where every attribute can uniquely distinguish every outcome. However, smaller values of k may overlook the ability of an attribute to distinguish winning agents from losing ones. Nonetheless, we observe that for all $k \geq 4$ (yielding 16 or more partitions) we can extract a consistent ranking.

Information gain example

To illustrate our use of information gain we will walk through the following short example. As in our primary analysis, we will consider the outcome value of a performance to be whether or not an agent placed first. In our example we will evaluate the information gained by knowing the maximum lead time an agent requested on any component order in the game. We will only consider the feature using 2 partitions: one for less than 50 and one for greater than or equal to 50. From 6 games we will create 36 performance observations (assuming 6 agents in each game). Note that $\frac{6}{36}$ are first place performances, giving the outcome variable over our dataset an entropy of:

$$H(O) = - \left[\frac{6}{36} \log \left(\frac{6}{36} \right) + \frac{30}{36} \log \left(\frac{30}{36} \right) \right] \approx 0.65$$

Now assume that 8 of the 36 performances had lead times greater than 50, including 5 of the 6 winning performances. In other words, the probability of observing a long lead time is $\frac{8}{36}$, the probability of an agent winning given that it had a long lead time is $\frac{5}{8}$, and the probability of observing a winning performance without a long lead time is $\frac{6-5}{36-8} = \frac{1}{28}$.

We can now calculate the conditional entropy of the outcome variable in the case where the maximum lead time attribute is greater than or equal to 50 (“long”) and when it is less than 50 (“short”),

$$\begin{aligned} H(O \mid \text{“long”}) &= - \left[\frac{5}{8} \log \left(\frac{5}{8} \right) + \frac{3}{8} \log \left(\frac{3}{8} \right) \right] \\ &\approx 0.95 \\ H(O \mid \text{“short”}) &= - \left[\frac{1}{28} \log \left(\frac{1}{28} \right) + \frac{27}{28} \log \left(\frac{27}{28} \right) \right] \\ &\approx 0.22 \end{aligned}$$

Using the conditional entropies we can calculate the average entropy of the outcome variable conditioned on the lead time attribute, A ,

$$\begin{aligned} H(O \mid A) &= P(\text{“long”})H(O \mid \text{“long”}) \\ &+ P(\text{“short”})H(O \mid \text{“short”}) \\ &\approx 0.38 \end{aligned}$$

Finally, the information gain of the outcome, O , from the attribute A , is the difference between the entropy of O independent of A , and its average entropy conditioned on A ,

$$\begin{aligned} IG(O, A) &= H(O) - H(O \mid A) \\ &\approx 0.27 \end{aligned}$$

Note that, because the initial entropy of the “first place” feature is about 0.65, the maximum possible information gain for any feature is also 0.65.

Information gain results

We now present the information gain of 20 different features across 6 values of k (representing 2, 4, 8, 16, 32, and 64 partitions). Our data set included all of the 80 games from the 2006 final rounds. Figure 3 shows the information gain of 6 different features at each level of discretization. It illustrates that upon reaching 16 or more partitions, features that provide more information tend to do so at finer discretization levels as well. Therefore, despite the potential drawbacks associated with the discretization process, we are still able to extract a fairly consistent ranking of features based on their ability to differentiate winning agents.

Figure 4 shows the information gain for all 20 different features ranked into 8 categories that are consistent from 16 partitions on. The ranking illustrates that the two features providing the most information about an agent’s performance were both related to its decisions about lead times on component orders. Additionally, 8 of the top 10 features in the ranking were related to decisions about component orders, such as their average quantity and reserve prices. Notably absent from the top distinguishing features were all demand-oriented features: the highest of these, total sell quantity (in revenue), tied with four other features for rank 7. This suggests that top agents were able to distinguish themselves primarily based on the collection of features that composed their procurement strategy (which is consistent with previous findings in (Benisch *et al.* 2006) regarding the 2005 seeding rounds).

³We later extend this technique to consider other outcomes: specifically, whether or not an agent finished in the top 3 positions.

⁴This is because performances in separate partitions remain separated as k increases.

Rank	Feature	2 partitions	4 partitions	8 partitions	16 partitions	32 partitions	64 partitions
1	Maximum lead time (supply)	0.217	0.257	0.277	0.347	0.385	0.418
2	Average lead time (supply)	0.225	0.245	0.310	0.339	0.368	0.404
3	Average early component order quantity (sent before day 25)	0.249	0.274	0.309	0.324	0.354	0.391
4	Average reserve price (supply)	0.218	0.257	0.286	0.304	0.314	0.345
	Small component order percentage (quantity ≤ 100)	0.211	0.212	0.239	0.286	0.316	0.341
	Average reserve price slack ^a (supply)	0.217	0.251	0.269	0.294	0.312	0.334
7	Last-minute component order percentage (lead time ≤ 3)	0.199	0.207	0.224	0.262	0.285	0.324
	Short lead time component order percentage (lead time ≤ 10)	0.220	0.227	0.259	0.277	0.289	0.309
	Total revenue (demand)	0.217	0.260	0.261	0.274	0.287	0.304
	Total quantity sold (demand)	0.212	0.245	0.259	0.262	0.275	0.302
11	Average quantity ordered per day (supply)	0.210	0.239	0.252	0.259	0.271	0.300
12	Average RFQ due date (demand)	0.201	0.209	0.232	0.243	0.265	0.291
13	Average factory utilization	0.201	0.204	0.209	0.222	0.230	0.274
	Average selling price (demand)	0.198	0.202	0.220	0.225	0.233	0.246
	Average purchase price (supply)	0.200	0.202	0.211	0.217	0.240	0.244
	Minimum bank account value	0.194	0.196	0.200	0.220	0.222	0.241
	Purchase price standard deviation (supply)	0.200	0.201	0.204	0.206	0.215	0.239
	Average stock value	0.195	0.200	0.201	0.223	0.229	0.238
	Average order price advantage (supply)	0.200	0.201	0.205	0.226	0.234	0.236
	Unsold stock at end of game	0.196	0.202	0.216	0.216	0.217	0.227

^aThe difference between reserve price specified and actual price paid.

Figure 4: The information gain of the 20 different features we tested at each level of discretization. The features are sorted by information gain at 64 partitions and ranked into groups that are distinguishable at each discretization level from 16 to 64 partitions.

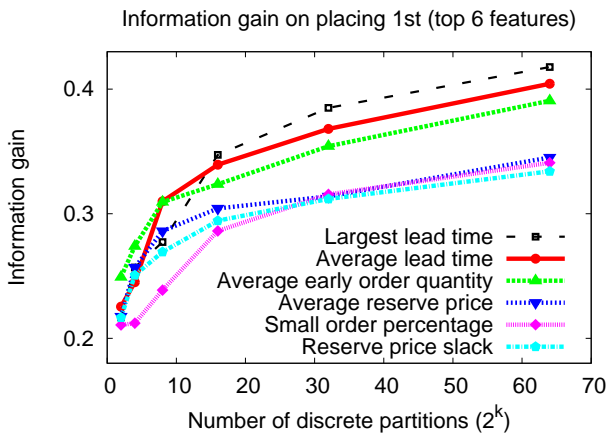


Figure 3: A plot showing the information gain for 6 different features at varying levels of discretization ($k \in \{1, \dots, 6\}$). The maximum possible information gain of any feature is ≈ 0.65 .

When calculating information gain for a feature, we determine the percentage of 1st place performances which occupy each partition for each feature, and likewise for the percentage of 2nd-6th place performances. Once we've identified an interesting feature, we can examine this information more directly with a histogram, showing us where exactly the distinctions between agents could be made. Figure 5, for example, shows a histogram comparing the percentage of 1st place performances in each of 16 partitions with the percentage of 2nd-6th place performances in those partitions for maximum component order lead times.

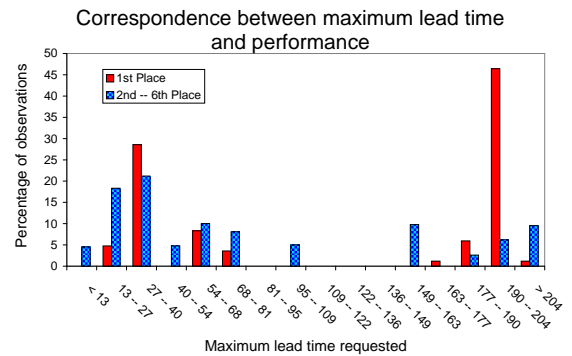


Figure 5: A histogram comparing the percentage of 1st place performances in each of 16 partitions with the percentage of 2nd-6th place performances in those partitions for maximum component order lead times. Our performance observations include all games from the 2006 final rounds.

We can see from this plot that a striking plurality of the winning performances used very long maximum lead times – from 190 to 204 days – while the second most prominent winning performance tended to keep maximum lead times at only 27 to 40 days. This clues us in to two strong strategies from the 2006 final rounds: winning agents tended to either order components almost to the end of the game at the very beginning, or they were more conservative and did not risk long lead times. Agents who restricted themselves to even shorter time ranges, or who took the large middle ground between 40 and 190 days, did not tend to be as successful.

Figure 6 shows a similar histogram examining the second most distinguishing feature: mean component order lead time. In this plot we see that, although a large maximum lead time was beneficial, agents who used long lead times excessively did not tend to perform well. Very few wins are observed for mean lead times greater than 40, while the plurality of lead times for winning performances sits at the relatively low range of 13 to 18. Finally we can see that for both wins and losses, the lower average lead times were a more popular choice.

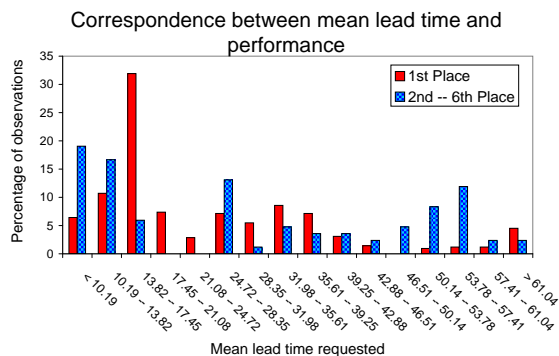


Figure 6: A histogram comparing the percentage of 1st place performances in each of 16 partitions with the percentage of 2nd-6th place performances in those partitions for average component order lead times. Our performance observations include all games from the 2006 final rounds.

Note that this analysis, by examining what distinguishes first place agents, focuses on a relatively small set of the agents, since many of the agents never, or rarely, placed first. For example, the very long maximum lead times which were strongly associated with first place performances were only used by 2 different agents. So while the results so far provide interesting clues about what may have set the few exceptionally successful agents apart from the rest, we also want to examine what more widely used behaviors were associated with success. To do so, we re-define our measure of success from “the agent placed first” to “the agent placed at least third.”

The results of this extension are shown in Figure 7 and a graphical version for the top 6 features is shown in Figure 8. These figures illustrate that the relative ordering of features is less consistent than before. Nonetheless, we are still able to extract 6 distinct levels of informativeness across 32 and 64 partitions. Many of our observations about first place agents hold true in this new ranking: decisions about component ordering continue to dominate the ranking, taking 9 of the 12 top spots. If we rank on the information gained at 64 partitions, all features in the top 10 previously remain in the top 10. There are certainly differences in the ranking – maximum lead time, for example, has fallen from being the most important feature to being third most important – but features which differentiated first place agents appear to continue to differentiate successful agents more generally.

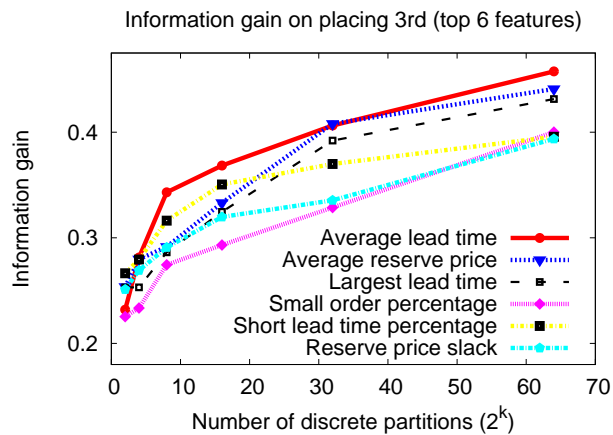


Figure 8: A plot showing the information gained about whether or not an agent placed 3rd or better for 6 different features at varying levels of discretization ($k \in \{1, \dots, 6\}$).

Discussion

This paper presented an investigation into which collection of behavioral features differentiated winning TAC SCM agents during the 2006 final rounds. We began with a visual inspection of games from one bracket of the 2006 semi-finals. Plots from these games revealed unique patterns, or “fingerprints,” which allowed us to isolate behavioral features that distinguished top performing agents in this bracket.

We then extended this analysis by applying a quantitative technique to all of the 80 games in the 2006 final rounds. This technique involved calculating the amount of information gained about an agent’s performance by knowing its value for each of 20 different features. The most informative features turned out to be related to direct decisions regarding component orders, such as the lead times and reserve prices used. These features differentiated winning agents in the 2006 final rounds significantly more than those related to costs and revenues.

Our information gain-based analysis technique was limited to examining the informativeness of individual features. Extending our technique to consider the effects of combinations of features may provide additional insight. For example, knowing an agent’s average selling price and average buying price together would probably be very informative. However, this raises additional concerns about over-fitting: using several features at once may uniquely identify each agent, instead of their shared characteristics.

As previously mentioned, our information gain-based technique can also be extended to consider other outcomes. For example, it may be interesting to investigate which features distinguish the *worst* agents. This can be accomplished by simply changing the outcome variable associated with each performance observation.

Finally, an agent designer may wish to answer the question, “what features differentiate games her agent wins from games it doesn’t?” This can be accomplished by modifying the information gain technique in the following ways. First,

Rank	Feature	2 partitions	4 partitions	8 partitions	16 partitions	32 partitions	64 partitions
1	Average lead time (supply)	0.232	0.281	0.343	0.368	0.406	0.458
	Average reserve price (supply)	0.254	0.280	0.291	0.333	0.408	0.441
3	Maximum lead time (supply)	0.253	0.253	0.286	0.325	0.392	0.431
4	Small component order percentage (quantity ≤ 100)	0.225	0.233	0.274	0.293	0.329	0.400
	Short lead time component order percentage (lead time ≤ 10)	0.266	0.279	0.316	0.350	0.370	0.396
	Average reserve price slack (supply)	0.251	0.269	0.291	0.320	0.336	0.394
	Last-minute component order percentage (lead time ≤ 3)	0.234	0.257	0.306	0.328	0.337	0.387
	Total revenue (demand)	0.246	0.303	0.310	0.350	0.358	0.379
	Average early component order quantity (sent before day 25)	0.283	0.309	0.321	0.330	0.339	0.377
	Total quantity sold (demand)	0.239	0.308	0.323	0.327	0.343	0.363
	Average quantity ordered per day (supply)	0.239	0.295	0.309	0.319	0.334	0.355
	Average order price advantage (supply)	0.225	0.231	0.244	0.318	0.334	0.340
13	Average order price (supply)	0.225	0.233	0.283	0.290	0.316	0.330
14	Average RFQ due date (demand)	0.215	0.217	0.233	0.243	0.268	0.299
	Average stock value	0.206	0.217	0.221	0.262	0.271	0.285
16	Average factory utilization	0.194	0.194	0.195	0.207	0.236	0.284
	Minimum bank account value	0.204	0.210	0.222	0.229	0.248	0.270
	Average selling price (demand)	0.220	0.228	0.235	0.243	0.249	0.260
	Unsold stock at end of game	0.211	0.233	0.243	0.248	0.252	0.260
	Purchase price standard deviation (supply)	0.225	0.226	0.236	0.240	0.240	0.254

Figure 7: The information gain of the 20 different features we tested at each level of discretization, with respect to 3rd-place or better performances. The features are sorted by information gain at 64 partitions and ranked into groups that are distinguishable with 32 and 64 partitions. The maximum possible information gain of any feature is 1.

only consider performance observations of the agent in question. Second, use features related to the game overall, such as its average customer demand, rather than features of a specific agent's behavior.

Acknowledgements

This work was supported in part by the National Science Foundation ITR grant 0205435, by a grant from SAP, and by Carnegie Mellon University's e-SCM Lab.

References

- Arunachalam, R., and Sadeh, N. 2005. The supply chain trading agent competition. *Electronic Commerce Research Applications* 4(1):63–81.
- Benisch, M.; Andrews, J.; Bangerter, D.; Kirchner, T.; Tsai, B.; and Sadeh, N. 2005. CMieux analysis and instrumentation toolkit for TAC SCM. Technical Report CMU-ISRI-05-127, School of Computer Science, Carnegie Mellon University.
- Benisch, M.; Andrews, J.; Sardinha, A.; and Sadeh, N. 2006. CMieux: Adaptive strategies for supply chain management. In *Proceedings of International Conference on Electronic Commerce (ICEC)*.
- Borghetti, B.; Sodomka, E.; Gini, M.; and Collins, J. 2006. A market-pressure-based performance evaluator for TAC-SCM. In *Proceedings of AAMAS Workshop on Trading Agent Design and Analysis (TADA)*.
- Chopra, S., and Meindl, P. 2004. *Supply Chain Management*. New Jersey: Pearson Prentice Hall.
- Collins, J.; Arunachalam, R.; Sadeh, N.; Eriksson, J.; Finne, N.; and Janson, S. 2006. The supply chain management game for 2006 trading agent competition (TAC SCM). Technical Report CMU-ISRI-05-132, School of Computer Science, Carnegie Mellon University.
- He, M.; Rogers, A.; David, E.; and Jennings, N. R. 2005. Designing and evaluating an adaptive trading agent for supply chain management applications. In *Proceedings of IJCAI Workshop on Trading Agent Design and Analysis (TADA)*.
- He, M.; Rogers, A.; Luo, X.; and Jennings, N. R. 2006. Designing a successful trading agent for supply chain management. In *Proceedings of Autonomous Agents and Multiagent Systems (AAMAS)*.
- Jordan, P. R.; Kiekintveld, C.; Miller, J.; and Wellman, M. P. 2006. Market efficiency, sales competition, and the bullwhip effect in the TAC SCM tournaments. In *Proceedings of AAMAS Workshop on Trading Agent Design and Analysis (TADA)*.
- Kiekintveld, C.; Vorobeychik, Y.; and Wellman, M. 2006. An analysis of the 2004 supply chain management trading agent competition. In Pour, H. L.; Sadeh, N.; and Janson, S., eds., *Agent-Mediated Electronic Commerce: Designing Trading Agents and Mechanisms*, number 3937 in Lecture Notes in AI, 99–112. Springer-Verlag.
- Mitchell, T. M. 1997. *Machine Learning*. McGraw Hill.
- Pardoe, D., and Stone, P. 2006. Predictive planning for supply chain management. In *Proceedings of Automated Planning and Scheduling*.
- Sadeh, N.; Hildum, D.; Kjenstad, D.; and Tseng, A. 1999. Mascot: an agent-based architecture for coordinated mixed-initiative supply chain planning and scheduling. In *Proceedings of Agents Workshop on Agent-Based Decision Support in Managing the Internet-Enabled Supply-Chain*.
- Wellman, M. P.; Jordan, P. R.; Kiekintveld, C.; Miller, J.; and Reeves, D. M. 2006. Empirical game-theoretic analysis of the TAC market games. In *Proceedings of AAMAS Workshop on Game-Theoretic and Decision-Theoretic Agents*.

Appendix

This appendix includes several graphs that were omitted from the main text due to space constraints.

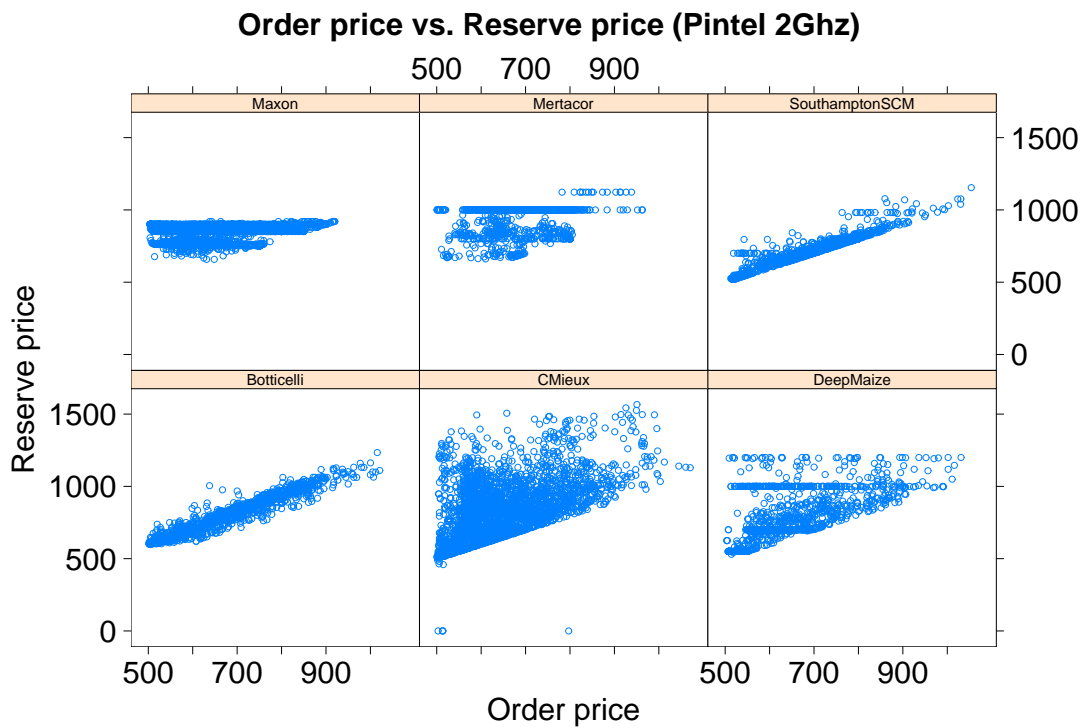


Figure 9: A plot showing the reserve price and order price of component orders placed by six different agents during the 2006 semi-finals.

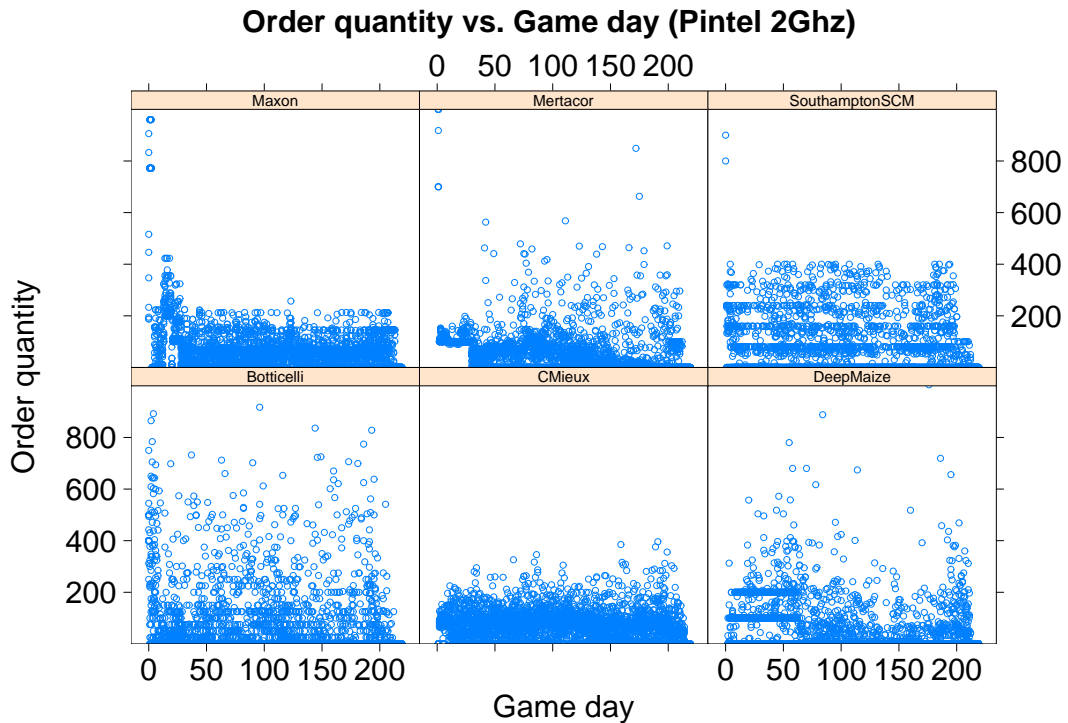


Figure 10: A plot showing the order quantity (clamped to 1000 to show detail) and game day of component orders placed by six different agents during the 2006 semi-finals.

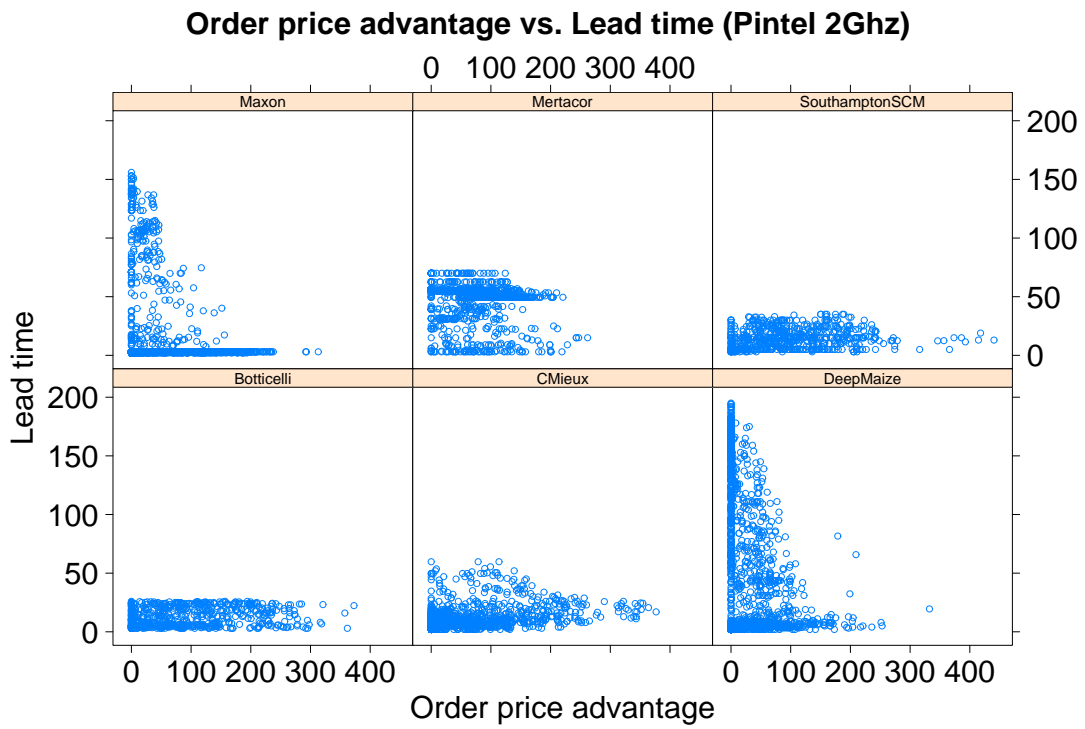


Figure 11: A plot showing the order price advantage and lead time of component orders placed by six different agents during the 2006 semi-finals.